6.1 Background: Consumer Values and Demand Curves

Let's consider a simplified relationship between price and quantity purchased by a single consumer, using some good, like a slice of pizza.

Suppose consumers are pretty much the same, and that they value the first slice at \$5, the second at \$4, the third at \$3, and so on. This is the marginal value. Knowing the value that consumers place on each subsequent slice allows us to construct Table 6.1, which shows the marginal value and total value for the various quantities. For the first slice, the total and marginal values are the same, both equal to \$5. For the second slice, the marginal value is \$4, while the total value of consuming two slices is \$9 = \$5 + \$4. For the third slice, the marginal value is \$3, and the total value is \$12 = \$5 + \$4 + \$3, and so on.

TABLE **6.1**
Pizza Value Table

Slices Purchased	Marginal Value (\$)	Total Value (\$)
1	5	5
2	4	9
3	3	12
4	2	14
5	1	15

The consumer uses marginal analysis to decide how much to consume because it is an extent decision. If the marginal value of consuming another unit is above the price, the consumer consumes another unit. For example, at a price of \$5, the consumer purchases only one slice because the second slice has a value (\$4) that is below the price.

At a price of \$4, the consumer purchases two slices; at a price of \$3, three slices; at a price of \$2, four slices; and at a price of \$1, five slices. The consumer's decision of how much to buy at each price is a demand curve, listed in Table 6.2.

TABLE **6.2**
Pizza Demand Schedule

Slice Price (\$)	Slices Purchased
5	1
4	2
3	3
2	4
1	5

A **demand curve** tells you how much consumers will purchase at a given price.

It is easy to see from Table 6.2 that the consumer purchases more as price falls, which is called the **first law of demand.** This makes intuitive sense. Consider the value you, a hungry consumer, receive from the first slice of pizza—it is likely to be substantial. The additional value you receive from eating the second slice is a bit less, and by the time you have eaten four slices, the additional value of the fifth is fairly small. The marginal, or additional, value of consuming each subsequent slice diminishes the more you consume.

In Table 6.3, we show how much surplus consumers get at each price. At a price of $5, there is no surplus as the consumer pays a price exactly equal to his total value. But as price declines, **consumer surplus**, the difference between total value and amount paid, increases. Note that the additional value of consuming the last slice is only $1.

TABLE 6.3
Pizza Consumer Surplus

Price ($)	Slices Purchased	Total Amount Paid ($)	Total Value ($)	Surplus ($)
5	1	5	5	0
4	2	8	9	1
3	3	9	12	3
2	4	8	14	6
1	5	5	15	10

To describe the buying behavior of a group of consumers, we add up all the individual demand curves to get an **aggregate demand curve.** The simplest way to show this is to consider the case where each consumer wants only a single item (i.e., the marginal value of a second unit is zero). To construct a demand curve that describes the behavior of seven buyers, we simply arrange the buyers by what they are willing to pay (e.g., $7, $6, $5, $4, $3, $2, and $1). At a price of $7, one buyer will purchase;[3] at a price of $6, two buyers will purchase; at $5, three buyers; and so on. At a price of $1, all seven buyers will purchase the good. An *aggregate or market demand curve* is the relationship between the price and the number of purchases made by this group of consumers. In Figure 6.1, we plot this demand curve.

Note that price—the independent variable—is on the "wrong" axis. There are good reasons for this that will become apparent, but for now, just accept that economists like to do things a little differently. Note also that economists have special jargon describing the response of demand to price. We say that as price decreases, "quantity demanded" increases. If something other than price causes an increase in demand, we instead say that the "demand shifts" to the right, or "demand increases," such that consumers purchase more at the same price. We'll discuss factors that shift demand in a later chapter.

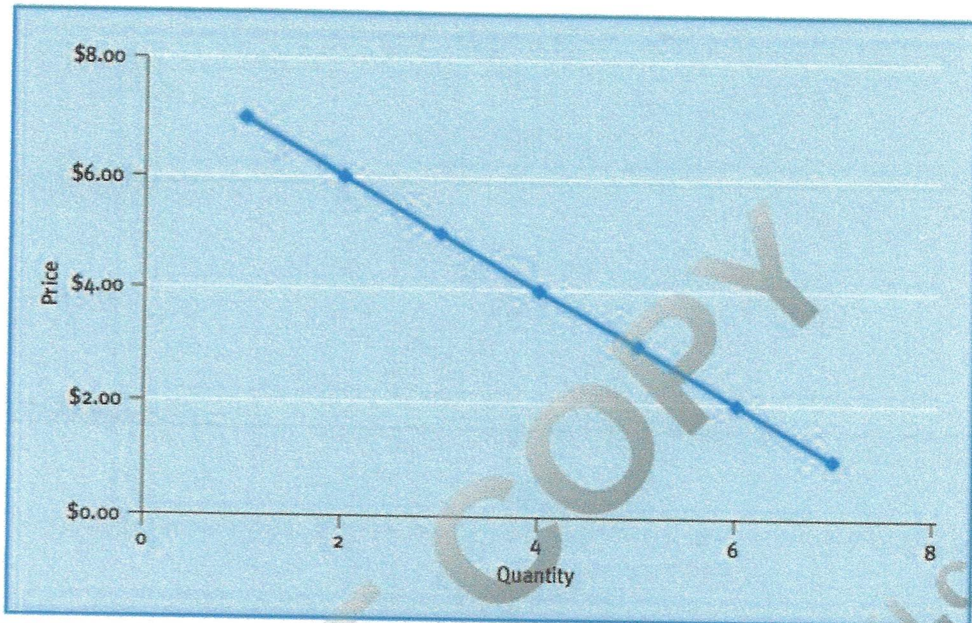

FIGURE 6.1 Demand Curve

To determine the quantity demanded at each price using the demand curve, look for the quantity on the horizontal axis corresponding to a price on the vertical axis. At a price of \$6, buyers demand two units; at a price of \$5, three units; and so on. As price falls, quantity demanded increases.

6.2 Marginal Analysis of Pricing

Demand curves present sellers with a dilemma. Sellers can raise price and sell fewer units, but earn more on each unit sold. Or they can reduce price and sell more, but earn less on each unit sold. This trade-off is at the heart of pricing decisions.

We use demand curves to change the pricing decision ("what price should I charge") into a quantity decision ("how much should I sell?") that we already know how to solve. If marginal revenue (MR) is greater than marginal cost (MC),[4] sell more, and you do this by reducing price. Reduce price (sell more) if MR > MC. Increase price (sell less) if MR < MC.

To see how to use marginal analysis to increase profit, examine Table 6.4. The columns list the Price, Quantity, Revenue, MR, MC, and total Profit for a simple demand curve. Suppose that the product costs \$1.50 to make. At a price of \$7, one consumer would purchase, so revenue would be \$7. Cost would be \$1.50, so profit on the first sale would be \$5.50.

TABLE 6.4
Optimal Price

Price ($)	Quantity	Revenue ($)	MR ($)	MC ($)	Profit ($)
7	1	7	7	1.50	5.50
6	2	12	5	1.50	9.00
5	3	15	3	1.50	10.50
4	4	16	1	1.50	10.00
3	5	15	−1	1.50	7.50
2	6	12	−3	1.50	3.00
1	7	7	−5	1.50	−3.50

If we sell a second unit, we have to reduce price to $6, and revenue increases from $7 to $12. We say that the MR of the second unit is $5. This is bigger than the MC of the second unit, so it pays to sell the second unit. In the last column, we see that profit increases from $5.50 to $9.00.

If we sell a third unit, we have to reduce price to $5, and revenue increases from $12 to $15. The MR of the third unit is $3, which is bigger than the MC of the third unit, so it pays to sell the third unit. In the last column, we see that profit increases from $9.00 to $10.50.

So far, all these changes have been profitable because MR has been greater than the MC. We earned $5.50 on the first unit, $3.50 on the second unit, and $1.50 on the third unit. These marginal profits sum to a total profit of $10.50, as indicated in the last column of Table 6.4.

However, if we sell a fourth unit, total profit decreases because the MR is only $1, which is less than the $1.50 MC. Selling three units, at a price of $5 is "optimal" because it makes the most profit.

After going through your analysis to compute the optimal price, suppose your boss looks at you and says, "This is the stupidest thing I've ever seen! Since the price is $5, and the cost of producing another good is only $1.50, we're leaving money on the table." What do you tell her?

Your boss has confused *average* revenue or price with *marginal* revenue. They're easy to confuse. Here's why. As long as price is greater than average cost, it appears that an increase in quantity would increase profit. However, this reasoning is incorrect because you cannot sell more without reducing the price on *all* goods, and not just on the extra units your boss wants to sell.

Tell your boss that you are already making all profitable sales—those for which MR exceeds MC. Marginal analysis, not average analysis, tells you where to price or, equivalently, how many units to sell.

6.3 Price Elasticity and Marginal Revenue

Unfortunately, you're never going to see a demand curve like the one in Figure 6.1. In general, it is very difficult to get information about demand at prices away from the current price. In fact, if anyone—particularly an economic consultant—ever tries to show you a complete demand curve, don't trust it; the consultant probably has only a very rough estimate as to what demand looks like away from current prices.

At this point you may be shaking your head and wondering why you have to learn about things you will never see. Table 6.4 shows us that we don't need the entire demand curve to know how to price—all we need is information on MR and MC. If MR > MC, reduce price; if MR < MC, increase price. As we saw earlier, marginal analysis points you in the right direction, but it doesn't tell you how far to go. You get to the best price by taking steps and then by re-computing MR and MC to see whether you should take another step.

So how do we estimate MR? The answer involves measuring quantity responses to past price changes, essentially "experimenting" with price changes, or surveying potential consumers to see how much they would buy in response to a price change. If you do get useful information about demand away from the current price, it's likely to be about the **price elasticity of demand**, which we denote by e.

$$e = \%\Delta \textit{Quantity Demanded} \div \%\Delta \textit{Price}$$

where %Δ means "percentage change in." Price elasticity measures how sensitive demand is to a change in price. A demand curve for which quantity changes more than price is said to be **elastic**, or sensitive to price; and a demand curve for which quantity changes less than price is said to be **inelastic**, or insensitive to price.

If $|e| > 1$, *demand is elastic; if* $|e| < 1$, *demand is inelastic.*

Since price and quantity move in opposite directions—as price goes up, quantity goes down, and vice versa—price elasticity is *always* negative (the first law of demand requires it); that is, $e < 0$. Because it is negative, many people will just drop the negative sign. To avoid confusion, we will usually refer to elasticity using its absolute value, denoted by $|e|$.

To show how to estimate elasticity, consider this 1999 "experiment" at MidSouth, a medium-sized retail grocery store. The store's managers decreased the price of three-liter Coke (diet, caffeine-free, and classic) from \$1.79 to \$1.50 because they wanted to match a price offered at a nearby Walmart. In response to the price drop, the quantity sold doubled, from 210 to 420 units per week.

To compute elasticity, simply take the percentage quantity increase and divide it by the percentage price decrease. Some confusion inevitably occurs because we can compute percentage changes in several different ways, depending on whether we divide the changes by their initial or final values. Usually, the best estimate comes from dividing by the midpoint of price $(P_1 + P_2)/2$ and the midpoint of quantity $(Q_1 + Q_2)/2$.[6]

$$e = [(Q_1 - Q_2) \div (Q_1 + Q_2)] \div [(P_1 - P_2) \div (P_1 + P_2)]$$

In the three-liter Coke example, the calculation works like this:

$$-3.8 = [(210 - 420) \div (210 + 420)] \div [(1.79 - 1.50) \div (1.79 + 1.50)]$$

In this case, the estimated price elasticity is −3.8, indicating an elastic demand, where a 1% decrease in price of three-liter Coke leads to a 3.8% increase in quantity.[7] The change in revenue associated with the change is

$$(\$1.50 \times 420) - (\$1.79 \times 210) = \$630 - \$375.90 = \$254.10$$

This experiment illustrates the relationship between elasticity and revenue,

for an elastic demand, a decrease in price leads to an increase in revenue.

In general, we can express the revenue change as the sum of quantity and price changes:[8]

$$\%\Delta Revenue \approx \%\Delta Quantity + \%\Delta Price$$

If quantity increases by more than price decreases (which it does for elastic demand), then revenue increases, and *vice versa*. The results in Tables 6.5 and 6.6 summarize the relationship.

TABLE 6.5
For *Elastic* Demand

Price increase → Revenue decrease
Price decrease → Revenue increase

TABLE 6.6
For *Inelastic* Demand

Price increase → Revenue increase
Price decrease → Revenue decrease

To illustrate the relationship between price, revenue, and elasticity, let's look at former mayor Marion Barry's tax increase on gasoline sales in the District of Columbia (DC). Before the tax was put into law, DC gas station owners argued against it, predicting that the 6% price increase would reduce quantity by 40%. Essentially, gas station owners were arguing that the price elasticity of demand for gasoline sold in the District was −6.7 = 40% ÷ 6%. Because of this very elastic demand, the gas station owners predicted that a tax increase would cause gasoline revenue, and the taxes collected out of revenue, to decline.

In fact, after the tax was levied, quantity fell by 38%, very close to what gas station owners had predicted. More importantly, tax revenue fell, as would be predicted by the top row of Table 6.5.

The exact numerical relationship between MR (change in revenue) and elasticity is MR = $P(1-1/|e|)$.[9] We can use this formula to express the marginal analysis rule using price elasticity and margins place of MR and MC:

$$MR > MC \textit{ implies that } (P - MC)/P > 1/|e|$$

This expression has an intuitive interpretation. The left side of the expresion is the *current margin* (P − MC)/P, whereas the right side is the *desired margin*, which is the inverse elasticity, 1/|e|. If the current margin is greater than the desired margin, reduce price because MR > MC. Intuitively, the more elastic demand becomes (1/|e| becomes smaller), the less you can profitably raise price because you will lose too many customers.

For example, after MidSouth Grocery reduced the price of three-liter Coke to $1.50, its actual margin was 2.7%, which is much less than the desired margin of 26%=1/|3.8|. In other words, the price was much too low according to our simple model of pricing (one product, one firm, one price). Ordinarily, a profit-maximizing store manager would raise the price in such a situation. In this case, however, the managers were using three-liter Coke as a *loss leader*, deliberately pricing too low as a way to attract customers to the store. Why? Because they hoped that customers would spend money on other items once they got there. Our simple pricing model does not take account of the effect of a low Coke price on sales of other products. We will discuss this, and other more complex pricing strategies, in later chapters.

To see how to use elasticity to set price, consider the following spreadsheet. The first two columns represent a demand curve, with price determining how much quantity is sold. In the third column, we compute the price elasticity. In the first row, at a price of $4.8, quantity is 7.0, and elasticity is −1.4. The actual margin is 0.46, and the desired margin is 0.71 = 1/(1.4). Since the actual margin is less than the desired margin, produce less, and you do this by increasing price.

In the second row, price goes up to $5.0, the actual margin increases to 0.48, and the desired margin decreases to 0.59 because demand becomes more elastic. At a price of 5.2, we see that profit is maximized at $16, and the desired margin (0.5) is equal to the actual margin. At the optimal price of 6.2 and a 50% margin, it is easy to compute MC (MC = 2.6).

Price ($)	Quantity	Elasticity	(P − MC)/P	1/|Elas|	Profit ($)
4.8	7.0	−1.4	0.46	0.71	15.50
5.0	6.6	−1.7	0.48	0.59	15.88
5.2	6.2	−2.0	0.50	0.50	16.00
5.4	5.7	−2.3	0.52	0.43	15.88
5.6	5.2	−2.7	0.54	0.37	15.52
5.8	4.7	−3.1	0.55	0.32	14.96
6.0	4.2	−3.5	0.57	0.29	14.22

6.4 What Makes Demand More Elastic?

Given the importance of price elasticity to pricing—the more elastic is demand, the lower is the profit-maximizing price—it's worthwhile to understand what makes demand more or less elastic. In this section, we list five factors that affect demand elasticity and optimal pricing.

Products with close substitutes have more elastic demand.

Consumers respond to a price increase by switching to their next-best alternative. If their next-best alternative is a very close substitute, then it doesn't take much of a price increase to induce them to switch. This is why revenues fell when Mayor Barry raised the price of gasoline by 6%. Since DC has many commuters, they began purchasing gasoline near their homes in Virginia and Maryland, close substitutes for gasoline in DC.

In a similar vein, we see that individual brands, such as Nike, have closer substitutes (other brands) than do aggregate product categories, like running shoes. This leads to our next factor.

Demand for an individual brand is more elastic than industry aggregate demand.

As a rough estimate, brand price elasticity is approximately equal to industry price elasticity divided by the brand share. For example, if the elasticity of demand for all running shoes is −0.4%, and the market share of Nike running shoes is 20%, price elasticity of demand for Nike running shoes is −0.4/(0.20) = −2. Using our optimal pricing formula, this would give Nike a desired margin of 50%.

If you search the Internet, you'll find industry price elasticity estimates that you can combine with market share estimates to get an estimate of brand elasticity. And you can use this estimate to gain a general idea of whether your brand price is too high or too low.

Products with many complements have less elastic demand.

Individual products that are consumed as part of a larger bundle of complementary goods—say, computers, operating systems, and the applications that run on them—have less elastic demand. One of the reasons that iPhones have such a low price elasticity of demand (and such a high margin) is due to the number of applications (apps) that run on them. If the price of an iPhone increases, you are less likely to substitute to another product, due to the complementary apps.

Another factor affecting elasticity is time. Given more time, consumers are more responsive to price changes. They have more time to find substitutes when price goes up and more time to find new uses for a good when price goes down. This leads to our fourth factor:

In the long run, demand curves become more elastic.

This phenomenon could also be explained by the speed at which price information is disseminated. As time passes, information about a new price becomes more widely known, so more consumers react to the change.

As an example, consider automatic teller machine (ATM) fees. In 1997, a bank in Evanston, Indiana, ran an experiment to determine the elasticity of demand for ATMs with respect to ATM fees. At a selected number of ATMs, the bank raised user fees from \$1.50 to \$2.00. When informed of the fee increase, users typically completed the current transaction (the short run) but avoided the higher-priced ATMs in the future (the long run).

Our final factor relates elasticity to the price level. As price increases, consumers find more alternatives to the good whose price has gone up. And with more substitutes, demand becomes more elastic.

As price increases, demand becomes more elastic.

For example, high-fructose corn syrup (HFCS) is a caloric sweetener used in soft drinks. For this application, sugar is a perfect substitute for HFCS. However, import quotas and price supports have raised the U.S. price of sugar to about twice that of HFCS. All soft drink bottlers have switched to HFCS from sugar. Because bottlers have no close substitutes for *low-priced* HFCS, its demand is less elastic. But if the price of HFCS were to rise to that of sugar, sugar would become a perfect substitute for HFCS. In other words, demand for *high-priced* HFCS becomes very elastic as it approaches the price of sugar.

As a strategy, many firms would like to make demand for their products less elastic (less sensitive to price). One way to do this is to create a brand identity that makes consumers less sensitive to price. A good example of this is Harley Davidson whose motorcycle brand has become synonymous with freedom, individualism, and the American Way. To build and maintain the brand the firm supports user groups (H.O.G., Harley Owners' Group) and sponsors motorcycle rallies (e.g., in Sturgis).

6.5 Forecasting Demand Using Elasticity

We can also use elasticity as a forecasting tool. With an elasticity and a percentage change in price, you can predict the corresponding change in quantity:[10]

$$\%\Delta Quantity \approx e(\%\Delta Price)$$

For example, if the price elasticity of demand is −2, and price goes up by 10%, then quantity is forecast to decrease by 20%.

Remember that price is only one of many factors that affect demand. Income, prices of substitutes and complements, advertising, and tastes all affect demand. To measure the effects of these other variables on demand, we define a factor elasticity of demand:

$$Factor\ elasticity\ of\ demand = (\%\Delta Quantity\ Demanded) \div \%\Delta Factor$$

Factors can be anything that affects demand, such as temperature, other prices, or incomes. For example, demand for bottled water, iced tea, and carbonated soft drinks is strongly influenced by temperature. If the temperature elasticity of demand for beverages is 0.25, then a 1% increase in temperature will lead to a 0.25% increase in quantity demanded.

Income elasticity of demand measures the change in demand arising from changes in income. Positive income elasticity means that the good is **normal**; that is, as income increases, demand increases. Negative income elasticity means that the good is **inferior**; that is, as income increases, demand declines. The decreasing incomes associated with the financial crisis of 2008 provided a number of examples of inferior goods. Although most retailers saw dramatic sales declines in 2008, Walmart's sales increased. Sales of Spam, a low-priced meat product, shot up in 2008, leading Hormel to add a second shift at its Minnesota factory.

Cross-price elasticity of demand for Good A with respect to the price of Good B measures the change in quantity demanded for A caused by a change in the price of B. A positive cross-price elasticity means that Good B is a **substitute** for Good A: as the price of a substitute increases, demand increases.

Negative cross-price elasticity means that Good B is a **complement** to Good A: as the price of a complement increases, demand decreases. Computers, for example, are complements to operating systems that run on them. We can trace part of Microsoft's initial success with its DOS operating system to its strategy of licensing DOS to competing computer manufacturers. That strategy helped keep the price of computers low, which stimulated demand for Microsoft's operating system.

We can estimate factor elasticities by using a formula analogous to the estimated price elasticity formula, and we can use factor elasticities to forecast or predict changes over time or even changes from one geographic area to another.

Suppose you're trying to compare the year-to-year performance of one of your regional salespeople over a period in which income grew by 3%. If demand for your products has an income elasticity of 2, you would expect quantity to increase by 6%. You don't want to reward the salesperson for increases in quantity that are largely unrelated to her effort. A performance measure more closely related to effort would subtract 6% from the actual growth because that is the growth related to income, and not to sales effort.

6.6 Stay-Even Analysis, Pricing, and Elasticity

Stay-even analysis is a simple two-step procedure that tells you whether a given price increase, for example, 5%, will be profitable.

1. In the first step, you compute how much quantity you can afford to lose before the price increase becomes unprofitable. This "stay-even quantity" is a simple function of the size of the price increase and the contribution margin, $\%\Delta Q = \%\Delta P/(\%\Delta P + \text{margin})$, where margin $= (P - MC)/P$.[11]
2. In the second step, you predict how much quantity will go down if you raise price by the given amount.

The decision rule is simple:

If the predicted quantity is less than the stay-even quantity, then the price increase will likely be profitable, and vice versa.

The analysis gives you a quick answer to the question of whether changing price is profitable. It uses the same information as does the marginal analysis of Section 6.4, but it does so in a simpler, intuitive way.